

The MSR System for Temporal Slot Filling at TAC 2013

Avirup Sil
Temple University, Philadelphia
avi@temple.edu

Silviu Cucerzan
Microsoft Research
silviu@microsoft.com

Outline

- Introduction to the Temporal Slot Filling Task
- Our Approach
 - Gathering Training Data from Wikipedia
 - Relationship Classifier
 - Date Classifier
- Experiments
- Conclusion and Future Work

Previous work on Relation Extraction

"Bill Clinton, the forty-second president of the US, was the first to pay down principle.."

- Output of Relation Extraction systems [Etzioni et. al, 05, Agichstein & Gravano, 00]:
 - President_of(Bill Clinton, United States)
- Limitation:
 - Does not capture temporal validity of the relationship
 - President_of(Bill Clinton, USA) is **true** during time-frame 1993-2001

The Temporal Slot Filling Task

■ Input:

- A binary relation
 - Example: `spouse (Brad Pitt, Jennifer Aniston)`
- A document supporting the relation

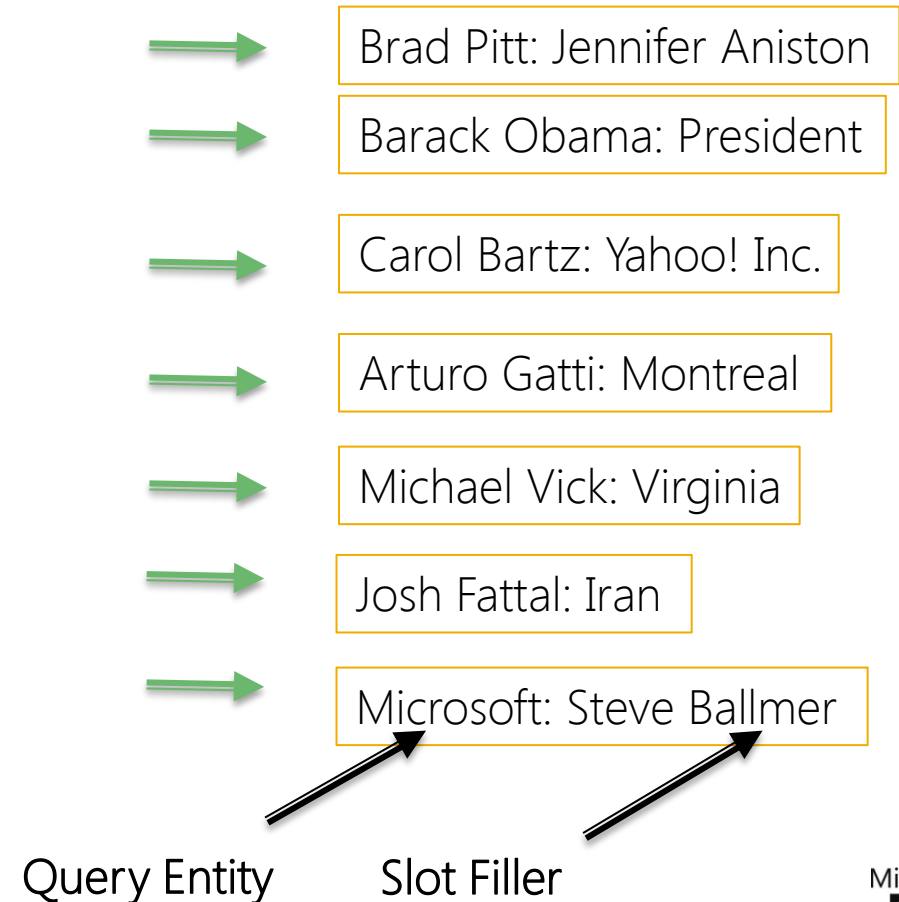
■ Output:

- A 4-tuple timestamp `[T1, T2, T3, T4]`
 - `[2000-07-29, nil, nil, 2005-10-02]`
- A sentence supporting the temporal validity of the relation
 - *"Pitt married Jennifer Aniston on July 29, 2000... the couple divorced five years later in October 2, 2005."*

7 Relation Types

- Text Analysis Conference (TAC): Temporal Slot Filling track has the following relation types:

1. Spouse
2. Title
3. Employee Of
4. Cities of Residence
5. States/Provinces of Residence
6. Countries of Residence
7. Top Employees/Members



Outline

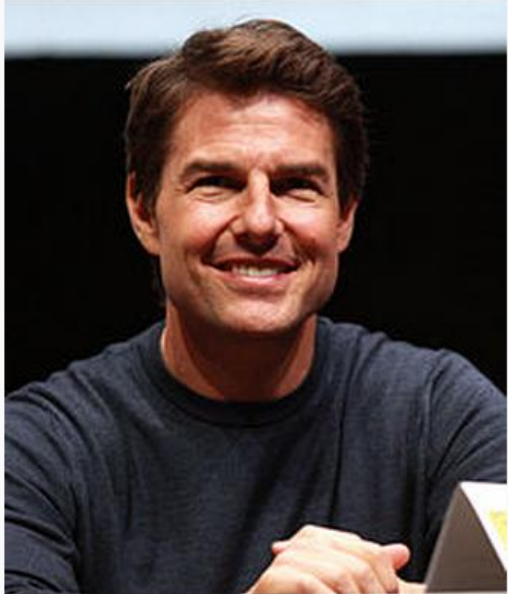
✓ Introduction to the Temporal Slot Filling Task

- Our Approach
 - Gathering Training Data from Wikipedia
 - Relationship Classifier
 - Date Classifier
- Experiments
- Conclusion and Future Work

Gathering Training Data using Distant Supervision

- No training data available
- We build our own training data from Wikipedia sentences
 - For every relation:
 - Extract Slot-Filler Names from Infoboxes from **all** Wikipedia pages
 - Apply MSR Entity Linker to resolve entity disambiguation and coreferences
 - Collect sets of contiguous sentences containing the slot-filler names
 - Build a language model by bootstrapping [Agichtein & Gravano, 00] textual patterns supporting the relations

Tom Cruise



Cruise at [San Diego Comic-Con International](#) in July 2013.

Born	Thomas Cruise Mapother IV July 3, 1962 (age 51) Syracuse, New York, U.S.
Occupation	Actor, producer, writer, director
Years active	1981–present
Religion	Scientology

Spouse: Katie Holmes

Gathering Training Data using Distant Supervision

- No training data available
- We build our own training data from Wikipedia sentences
 - For every relation:
 - Extract Slot-Filler Names from Infoboxes from **all** Wikipedia pages
 - Apply MSR Entity Linker to resolve entity disambiguation and coreferences
 - Collect sets of contiguous sentences containing the slot-filler names
 - Build a language model by bootstrapping [Agichtein & Gravano, 00] textual patterns supporting the relations

Wikipedia Sentences:

On October 6, 2005, Cruise and Holmes announced they were expecting a child..

.. On November 18, 2006, Holmes and Cruise were married at the 15th-century Odescalchi Castle in Bracciano, Italy...

On June 29, 2012, it was announced that Holmes had filed for divorce from Cruise after five and a half years of marriage.



Tom Cruise

From Wikipedia, the free encyclopedia



Katie Holmes

From Wikipedia, the free encyclopedia

Microsoft

Research

Gathering Training Data using Distant Supervision

- No training data available
- We build our own training data from Wikipedia sentences
 - For every relation:
 - Extract Slot-Filler Names from Infoboxes from **all** Wikipedia pages
 - Apply MSR Entity Linker to resolve entity disambiguation and coreferences
 - Collect sets of contiguous sentences containing the slot-filler names
 - Build a language model by bootstrapping [Agichtein & Gravano, 00] textual patterns supporting the relations

Patterns Extracted:

- DATE: X and Y were expecting a child
- DATE: X and Y were married
- DATE: X had filed for divorce from Y
- ...

X==Query Entity

Y== Slot Filler

We extract up to 5-grams.

Normalizing Date Surface Forms

- We run Stanford SUTime [Chang & Manning, 12] to resolve date surface forms

Raw Input Document:

```
<DOC id="AFP_ENG_20090626.0737" type="story" >
<HEADLINE>Distraught Madonna 'can't stop crying' over Jackson</HEADLINE>
<DATELINE>Los Angeles, June 25, 2009 (AFP)</DATELINE>
<TEXT><P>Pop diva Madonna revealed she was left in tears over the death of
Michael Jackson on Thursday, saying the music world had lost ..</P>
</TEXT>
</DOC>
```

Document normalized with Timestamps:

```
<DOC id="AFP_ENG_20090626.0737" type="story" >
<HEADLINE>Distraught Madonna 'can't stop crying' over Jackson</HEADLINE>
<DATELINE>Los Angeles, June 25, 2009 (AFP)</DATELINE>
<TEXT><P>Pop diva Madonna revealed she was left in tears over the death of
<TIMEX3 t0="2009-06-25">Thursday</TIMEX3> the music world had lost ..</P>
</TEXT>
</DOC>
```

RELCL: Identifying positive vs. Negative Sentences

■ Training:

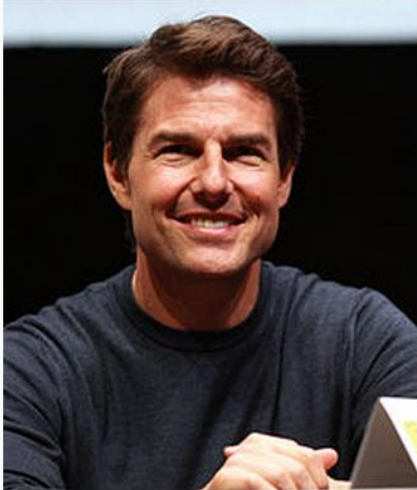
■ Example:

- Query Entity (X): Tom Cruise; Slot Filler (Y): Katie Holmes
- Sentence 1: *"On November 18, 2006, Holmes and Cruise were married in Bracciano, Italy.."*
- Sentence 2: *"In 2003, Cruise starred in the historical drama The Last Samurai.."*

Features	X and Y were married	Y, who died in DATE	were married in LOC	..	X married in DATE	X's wife Y	Y, who died	married	Label
Sentence 1	1	0	1	..	0	0	0	1	+1
Sentence 2	0	0	0	..	0	0	0	0	-1

■ Classifier:

- Boosted Decision Trees [Burges, 2010]

Tom Cruise

Cruise at San Diego Comic-Con International in July 2013.

Born Thomas Cruise Mapother IV
July 3, 1962 (age 51)
Syracuse, New York, U.S.

Occupation Actor, producer, writer, director

Years active 1981–present

Religion Scientology

Spouse: Katie Holmes

RELCL: Identifying positive vs. Negative Sentences

■ Testing:

■ Example:

- Query Entity: Norris Church
- Slot Filler: Norman Mailer

TAC TSF Eval Document

```
<DOC id="NYT_ENG_20101121.0120" type="story" >  
<HEADLINE>NORRIS CHURCH MAILER, ARTIST AND  
WRITER, DIES AT 61</HEADLINE>  
<TEXT>  
<P>Norman Mailer, whom Norris married in  
1980, was an attentive father..</P>  
<P>Norman Mailer, who died in 2007 at  
84, who dreamed up Church because he..</P>  
<P>Norris gave birth to John Buffalo in 1978  
and spent..</P>
```

RELCL: Identifying positive vs. Negative Sentences

■ Testing:

■ Example:

- Query Entity: Norris Church
- Slot Filler: Norman Mailer

TAC TSF Eval Document

```
<DOC id="NYT_ENG_20101121.0120" type="story" >
<HEADLINE>NORRIS CHURCH MAILER, ARTIST AND
WRITER, DIES AT 61</HEADLINE>
<TEXT>
<P> Y, whom X married in _DATE, was an attentive
father..</P>
<P> Y, who died in _DATE at 84, who dreamed
up X because he..</P>
<P> X gave birth to John Buffalo in _DATE
and spent..</P>
```

RELCL: Identifying positive vs. Negative Sentences

■ Testing:

■ Example:

- Query Entity: Norris Church
- Slot Filler: Norman Mailer

TAC TSF Eval Document

```
<DOC id="NYT_ENG_20101121.0120" type="story" >
<HEADLINE>NORRIS CHURCH MAILER, ARTIST AND
WRITER, DIES AT 61</HEADLINE>
<TEXT>
<P> Y, whom X married in _DATE, was an attentive
father..</P>
<P> Y, who died in _DATE at 84, who dreamed
up X because he..</P>
<P> X gave birth to John Buffalo in _DATE
and spent..</P>
```

Features	X and Y were married	Y, who died in DATE	were married in LOC	..	X married in DATE	X's wife Y	Y, who died	married
Sentence 1	0	0	0	..	1	0	0	1
Sentence 2	0	1	0	..	0	0	1	0
Sentence 3	0	0	0	..	0	0	0	0



RELCL: Identifying positive vs. Negative Sentences

- Testing:
 - Example:
 - Query Entity: Norris Church
 - Slot Filler: Norman Mailer

TAC TSF Eval Document

```
<DOC id="NYT_ENG_20101121.0120" type="story" >
<HEADLINE>NORRIS CHURCH MAILER, ARTIST AND
WRITER, DIES AT 61</HEADLINE>
<TEXT>
<P> Y, whom X married in _DATE, was an attentive
father..</P>
<P> Y, who died in _DATE at 84, who dreamed
up X because he..</P>
<P> X gave birth to John Buffalo in _DATE
and spent..</P>
```

Features	X and Y were married	Y, who died in DATE	were married in LOC	..	X married in DATE	X's wife Y	Y, who died	married
Sentence 1	0	0	0	..	1	0	0	1
Sentence 2	0	1	0	..	0	0	1	0
Sentence 3	0	0	0	..	0	0	0	0



RELCL: Identifying positive vs. Negative Sentences

■ Testing:

■ Example:

- Query Entity: Norris Church
- Slot Filler: Norman Mailer

TAC TSF Eval Document

```
<DOC id="NYT_ENG_20101121.0120" type="story" >
<HEADLINE>NORRIS CHURCH MAILER, ARTIST AND
WRITER, DIES AT 61</HEADLINE>
<TEXT>
<P> Y, whom X married in _DATE, was an attentive
father..</P>
<P> Y, who died in _DATE at 84, who dreamed
up X because he..</P>
<P> X gave birth to John Buffalo in _DATE
and spent..</P>
```

Features	X and Y were married	Y, who died in DATE	were married in LOC	..	X married in DATE	X's wife Y	Y, who died	married
Sentence 1	0	0	0	..	1	0	0	1
Sentence 2	0	1	0	..	0	0	1	0
Sentence 3	0	0	0	..	0	0	0	0



RELCL: Identifying positive vs. Negative Sentences

■ Testing:

■ Example:

- Query Entity: Norris Church
- Slot Filler: Norman Mailer

TAC TSF Eval Document

```
<DOC id="NYT_ENG_20101121.0120" type="story" >
<HEADLINE>NORRIS CHURCH MAILER, ARTIST AND
WRITER, DIES AT 61</HEADLINE>
<TEXT>
<P> Y, whom X married in _DATE, was an attentive
father..</P>
<P> Y, who died in _DATE at 84, who dreamed
up X because he..</P>
<P> X gave birth to John Buffalo in _DATE
and spent..</P>
```

Features	X and Y were married	Y, who died in DATE	were married in LOC	..	X married in DATE	X's wife Y	Y, who died	married
Sentence 1	0	0	0	..	1	0	0	1
Sentence 2	0	1	0	..	0	0	1	0
Sentence 3	0	0	0	..	0	0	0	0



D_{ATE}CL: Date Classification using Trigger Words

- **Goal**: Predict 4-tuple timestamp [T1, T2, T3, T4]
- **D_{ATE}CL**: A classifier using language models for “Start”, “End” and “In” predictors of relationship
 - Start predicts T1, T2; End predicts T3, T4; In predicts T2, T3
 - These compose of “Trigger Words”. Example for spouse relation:
 - Start: {married since _DATE, married SLOT_FILLER on,..}
 - End: {estranged husband QUERY_ENTITY, split in _DATE, SLOT_FILLER died,..}
 - In: {happily married, QUERY_ENTITY with his wife,..}

Extracting Timestamps

■ Example:

■ How to identify START?

- "Norman Mailer, whom Norris married in 1980, was an attentive father."
- "Y, whom X married in _DATE, was an attentive father."
- Indicates START of a "marriage" relationship
 - T1 = 1980-01-01; T2 = 1980-12-31; Justification_String: "1980"

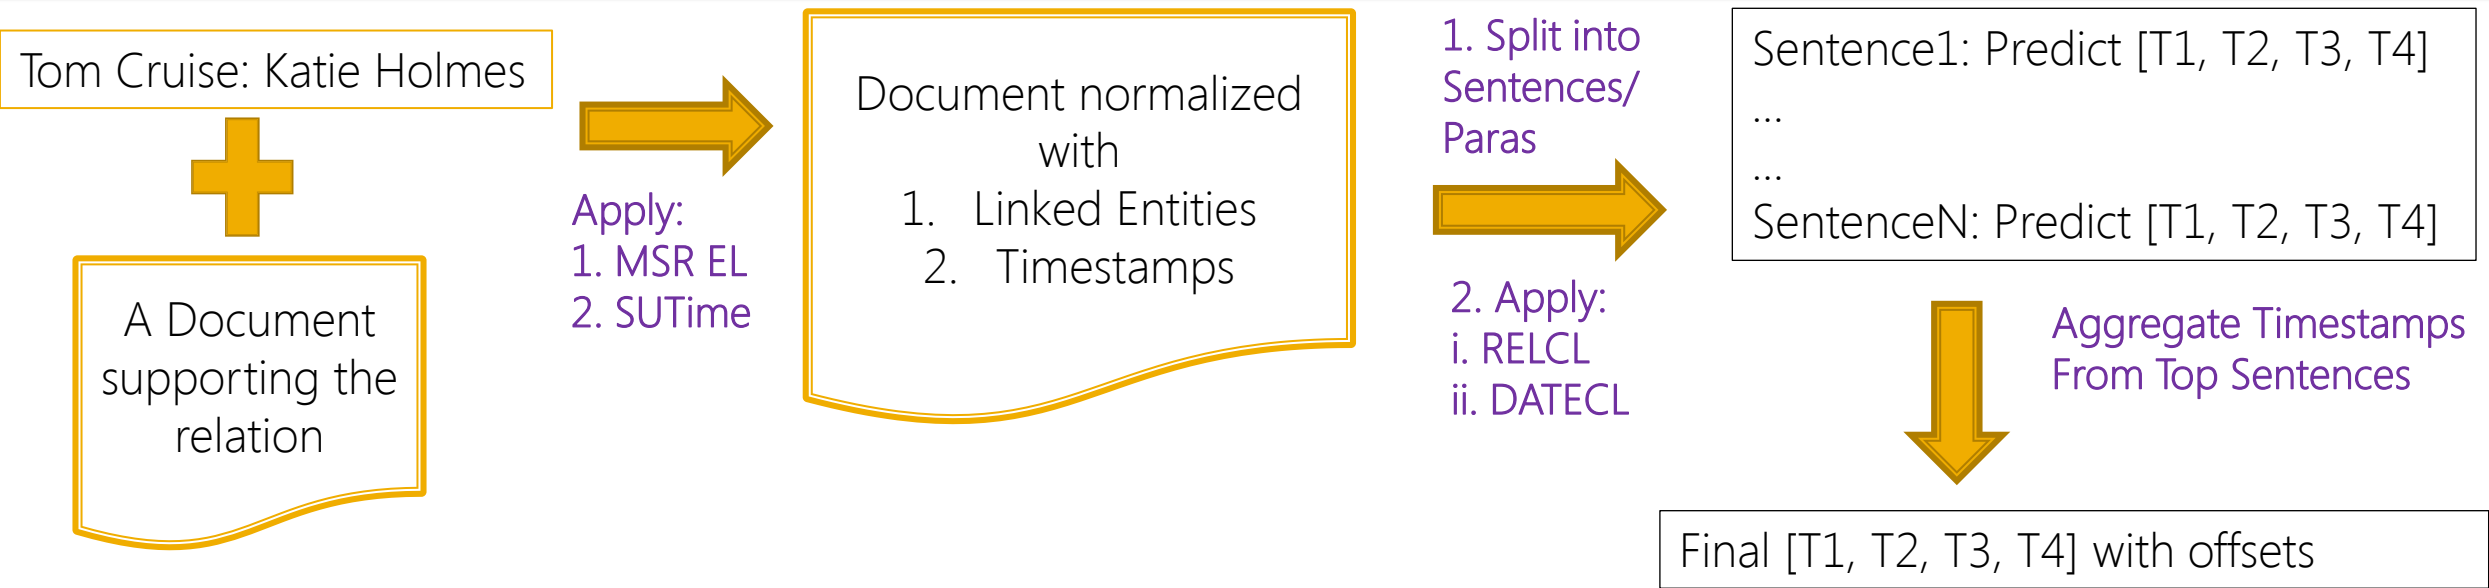
■ How to identify END?

- "Norman Mailer, who died in 2007 at 84,.."
- "Y, who died in _DATE at 84,.."
 - Indicates END of a "marriage" relationship
 - T3= 2007-01-01; T2 = 2007-12-31; Justification_String: "2007"

■ Aggregate the timestamps (based on Classifier confidence and heuristics):

- [1980-01-01 1980-12-31 2007-01-01 2007-12-31]

Predicting the timestamps (single-document prediction)



Update of the dates:

1. Initialize $T = [-\infty, +\infty, -\infty, +\infty]$
2. Iterate through classified timestamps
3. For a new T' aggregate :
 - $T \& T' = [\max(t_1, t_1'), \min(t_2, t_2'), \max(t_3, t_3'), \min(t_4, t_4')]$
 - Update only if $t_1 \leq t_2$; $t_3 \leq t_4$; $t_1 \leq t_4$

Outline

✓ Introduction to the Temporal Slot Filling Task

✓ Our Approach

- Gathering Training Data from Wikipedia
- Relationship Classifier
- Date Classifier

■ Experiments

■ Conclusion and Future Work

Experiments

- Dataset:
 - Wikipedia (May 2013)
 - Divide into Train and Dev
 - Train our RELCL and DATECL on Wikipedia training data
 - TAC
 - Training Data (7 examples; 1 per relation)
 - Evaluation Data (only for final test)
 - 273 examples (39 examples per relation)
- Evaluation Metric (as per TAC):
 - $S(\text{relation}) = \frac{1}{4} \sum_{i=1}^4 \frac{1}{1+d_i}$, $d_i = |r_i - k_i|$

Results

- On TAC 2013 Dataset

	Relations							
Run ID	Title	Spouse	EmployeeOf	CitiesOfRes	StatesOfRes	CountriesOfRes	Top_Employee	Overall
MS_MLI1	0.251	0.238	0.301	0.249	0.319	0.228	0.281	0.267
MS_MLI2	0.273	0.330	0.401	0.361	0.319	0.328	0.319	0.331

- Comparison:

Team	Mean Temporal Score (201 queries)
LDC	0.688 (Human)
MSR_TSF (Our System)	0.331
Team2	0.234
Team3	0.148
Team4	0.115
Team5	0.051

Conclusion and Future Work

- Wikipedia data proved to be an effective resource for the TSF task
 - Best performance in the task
- In the absence of annotated data **distant supervision** becomes effective
- Future (and ongoing) Work:
 - Using more than 1 single document for extracting Timestamps
 - Perform Joint-Relation extraction and Temporal Constraint attachment

Thanks!

- per: internOf(Avirup Sil, MSR): [2013-06-10,--,--,2013-09-06]
- Email:
 - avi@temple.edu
 - silviu@microsoft.com